High Performance Computing in the U.S. - the Next Five Years

Horst D. Simon

NERSC, Lawrence Berkeley Natl. Lab.

simon@nersc.gov

January 30, 1998

What is NERSC?

- National Energy Scientific Computing Center located at LBNL in Berkeley, California
- Together with ESnet the largest unclassified combination of supercomputing and networking power in the U.S.
- Supports DOE energy research programs
- Thousands of users nationally

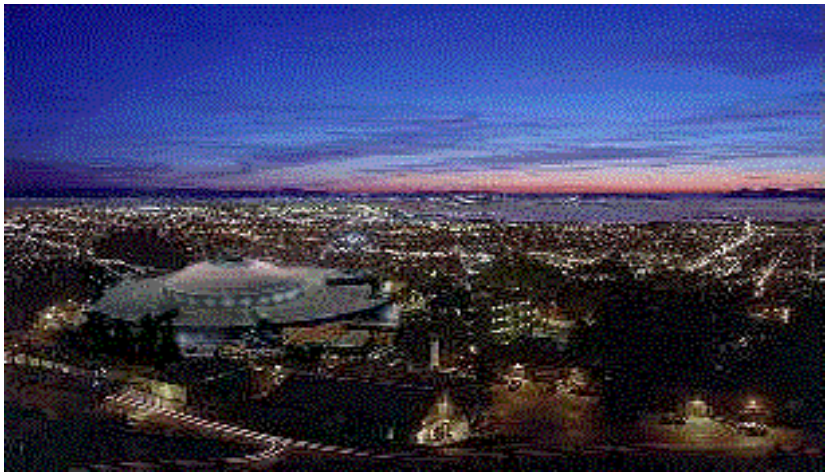

What is NERSC? (con't.)

- Relocated in 1996 from Livermore to Berkeley
- Close collaboration with UC Berkeley Computer Science
- Integration of Computer Science into high performance computing
- Grand Challenge Science Applications

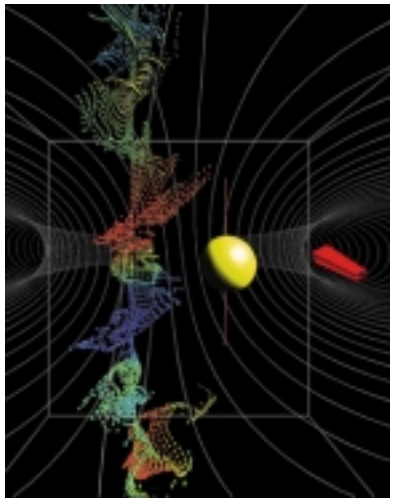

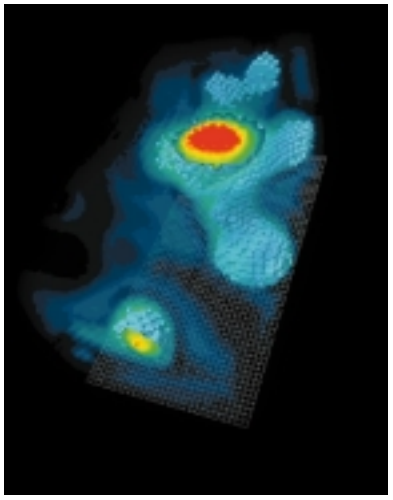

"Technology does not drive change at all. Technology merely enables change. It's our collective cultural response to the options and opportunities presented by technology that drives change."

Paul Saffo
Institute for the Future
Menlo Park, California

Overview

1992-1997 1998 1998-2002

technology

markets/business

government

non-linear events

Things that did not happen in the last five years

- TV and PC converge
- interactive TV
- video servers instead of video stores
- Apple/IBM/Motorola
- Intel makes a mistake
- MPPs go mainstream

after Forest Baskett, SGI/Cray

1992 - 1997: Technology

Moore's First Law

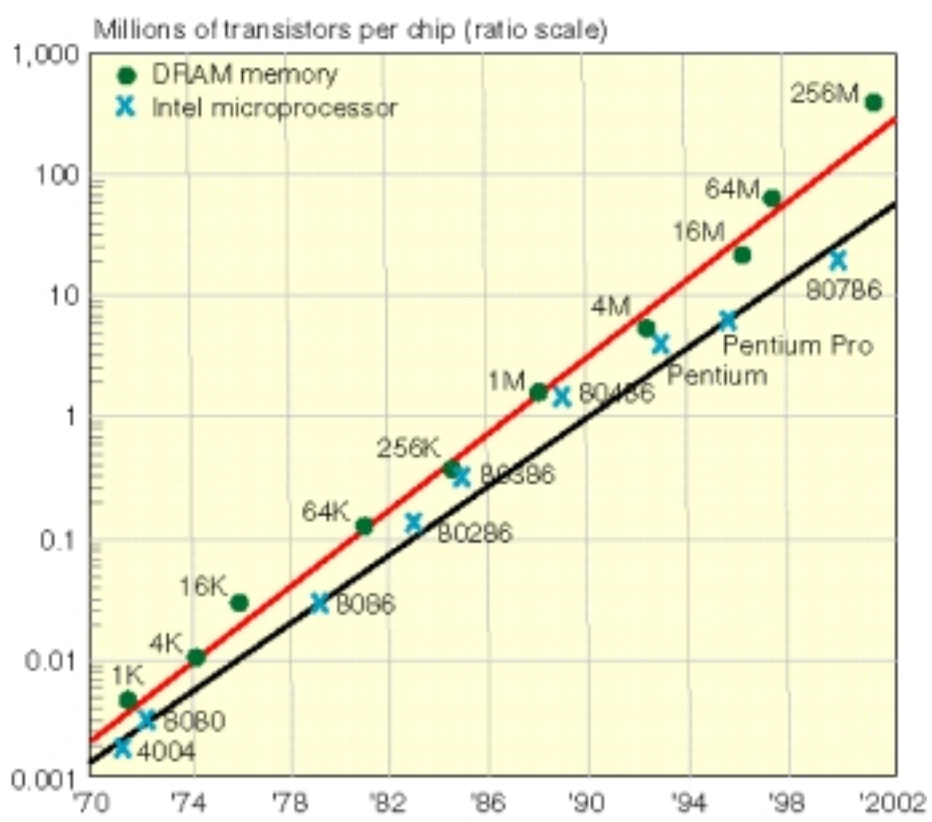

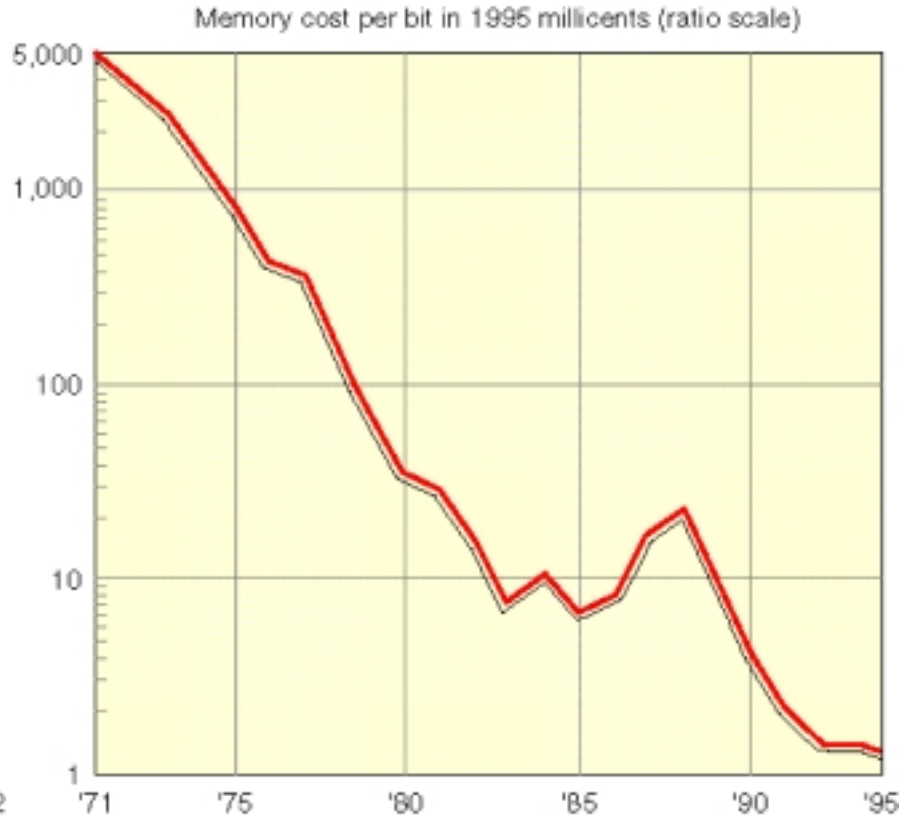

Source: VLSI Research Inc.

1992-1997: Technology

In the 1980's there have been fundamental changes in the microprocessor development ("killer micros")

- **— dramatic increase in number of transistors available per chip**
- **— architectural advances including the use of RISC ideas, pipelining and caches**
- **— as a result CPU performance has improved by a factor of 1.5 to 2.0 per year**

Maturation in the late 80s

Full impact in the early 90s

Microprocessors vs. Vector Supercomputers ca. 1994

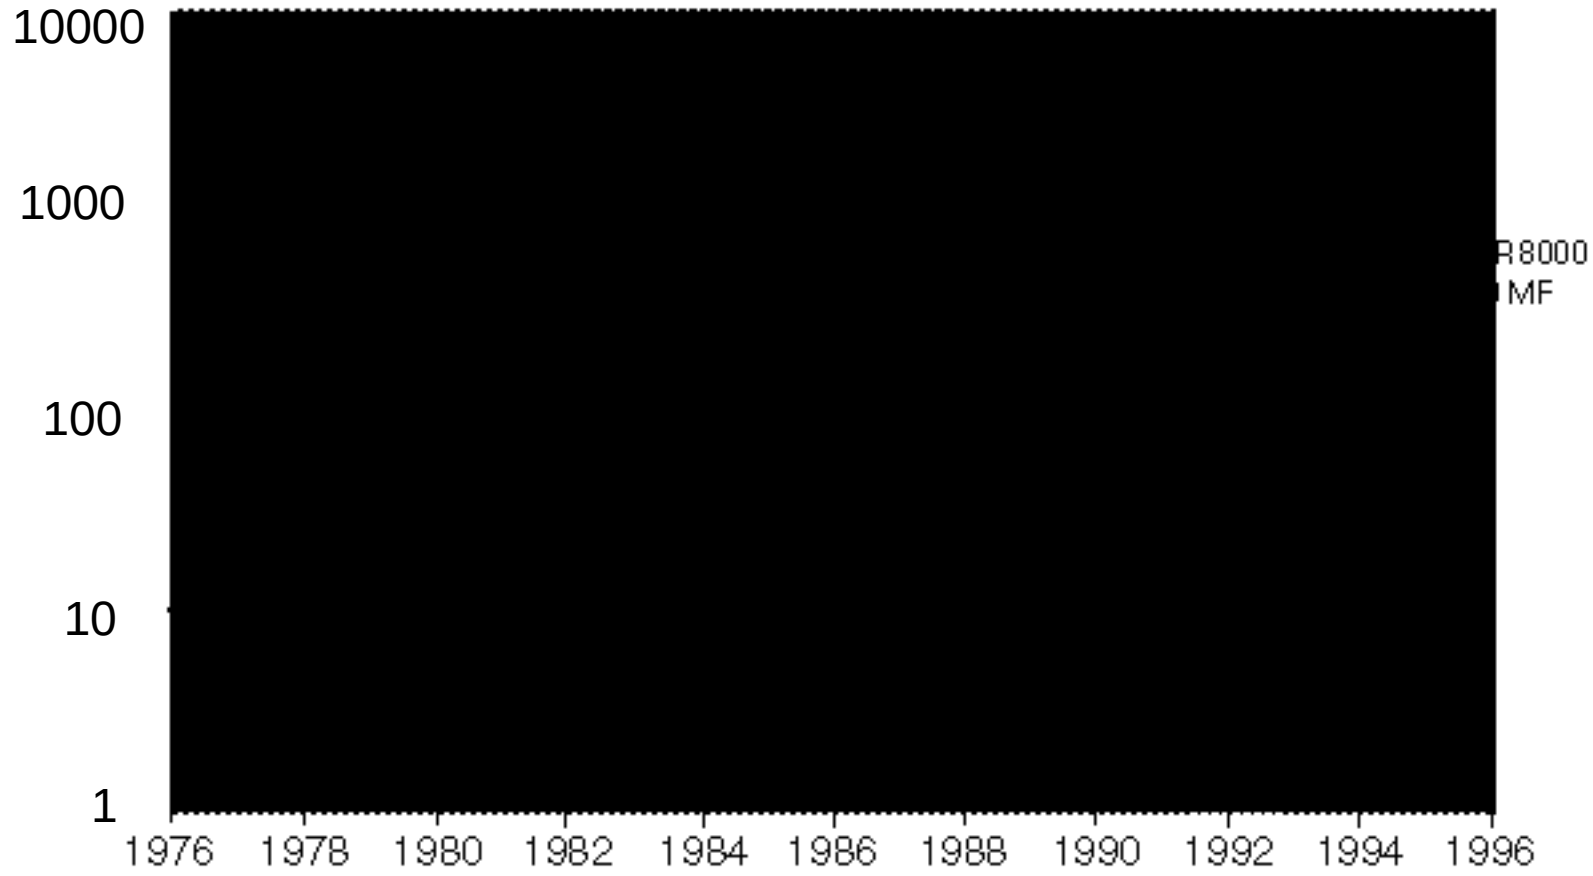

Top 500 - CPU Technology

Top 500 -total performance increase

Source http://parallel.rz.uni-mannheim.de/top500.html

1992-1997: HPC Market

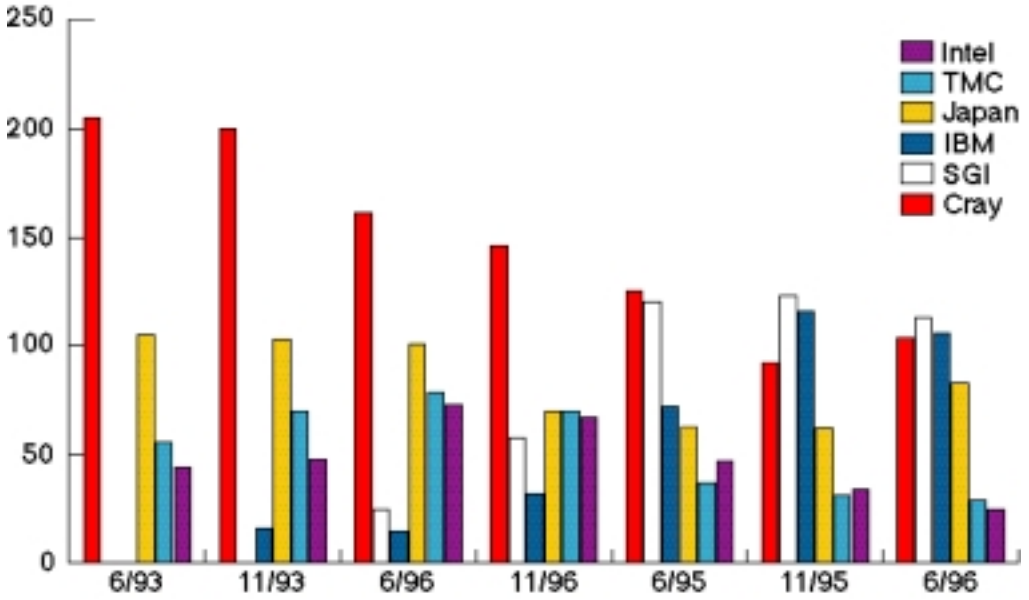

Data	Cray	SGI	IBM	Japan	TMC	Intel
6/93	205	0	0	105	56	44
11/93	200	0	16	103	70	48
6/96	161	25	15	101	79	73
11/96	146	58	32	70	70	67
6/95	125	120	72	63	37	47
11/95	92	123	116	62	31	34
6/96	104	113	106	83	29	25

The Revolution of 1994 - Major HPC Market Realignment

1991 Newcomers with CMOS and MPP technology (Intel, TMC, KSR) gain mind share and market share

1993 Cray, IBM, Convex go CMOS (T3D, SP 1/2, SPP 1000)

1994 TMC, KSR go out of business; SGI's SMP success

1995 HP buys Convex; Fujitsu, NEC introduce CMOS vector machines

1996 SGI buys Cray

1997 TOP500 take over by CMOS complete

1992-1997: HPC Market

1992-1997: Government

The years of the HPCC initiative.

The 1992 strategic goals were:

- **extend US technological leadership in HPC**
- **provide wide dissemination of HPC technology to serve national economy, national security, education, and environment**
- **spur gains in US productivity**

Coordination of HPC in government at highest level (NCO)

Spending increase (92 vs 91) of 30%,

thereafter about $800M/year

1997: The HPCC Impact

LAWRENCE BERKELEY NATIONAL LABORATORY

HPCC accomplished all of its strategic goals.

- **US leadership firmly established**
- **HPC has made firm inroads into the commerical world (MPP for database)**
- **new scientific culture has been established (large scale collaborations in computational science; new generation of computationally minded scientists)**

1997: The New HPC Marketplace

All majorUS HPC companies are now vertically integrated (SGI/Cray, IBM, HP, Sun DEC).

Almost all high-end procudcts are based onworkstation technology.

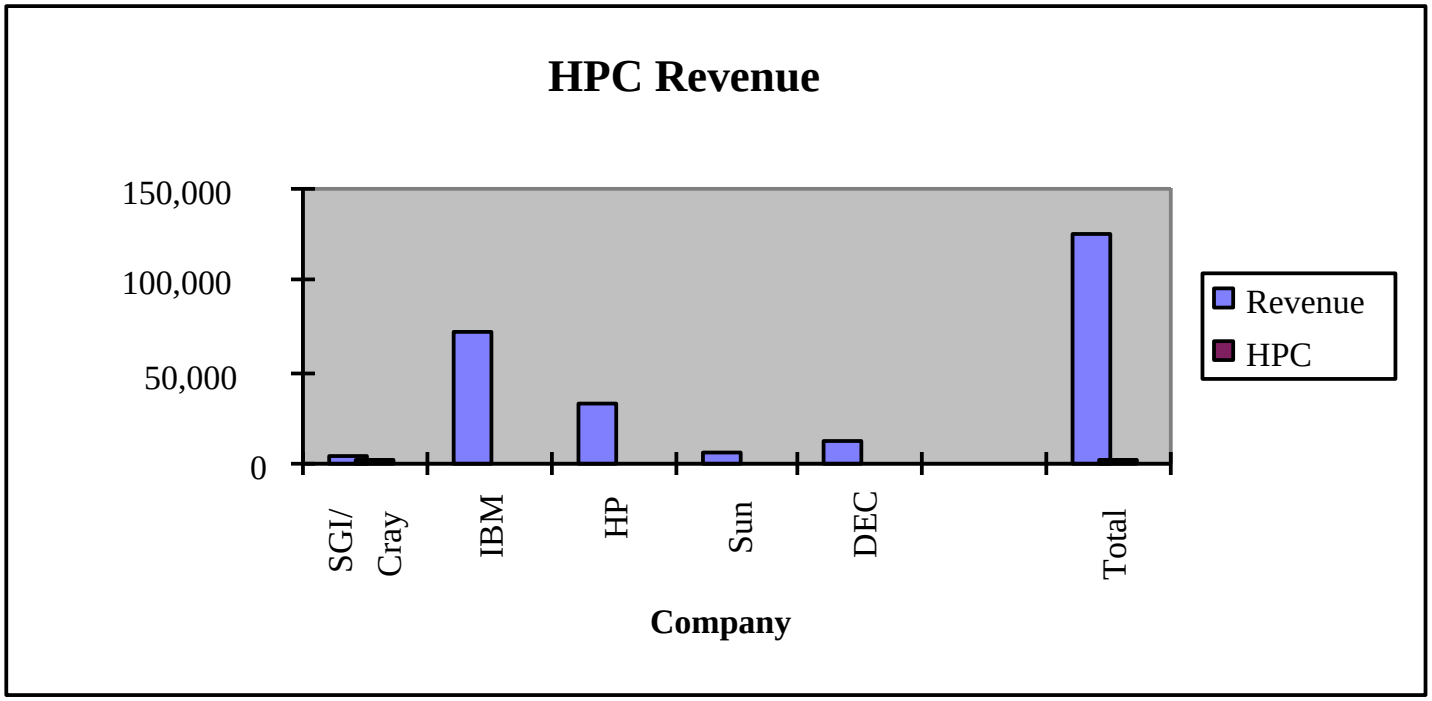

1997: The New HPC Marketplace

All these companies are in the computer business.

HPC customers must get used to a new role: they are no longer the center of attention.

Companies must have committment to technology, and understand the potential of technology leverage from the high-end, in order to remain in the HPC business.

Fortunately for us, the HPC users, they all do understand that (for now).

1997: The HPC Business Model

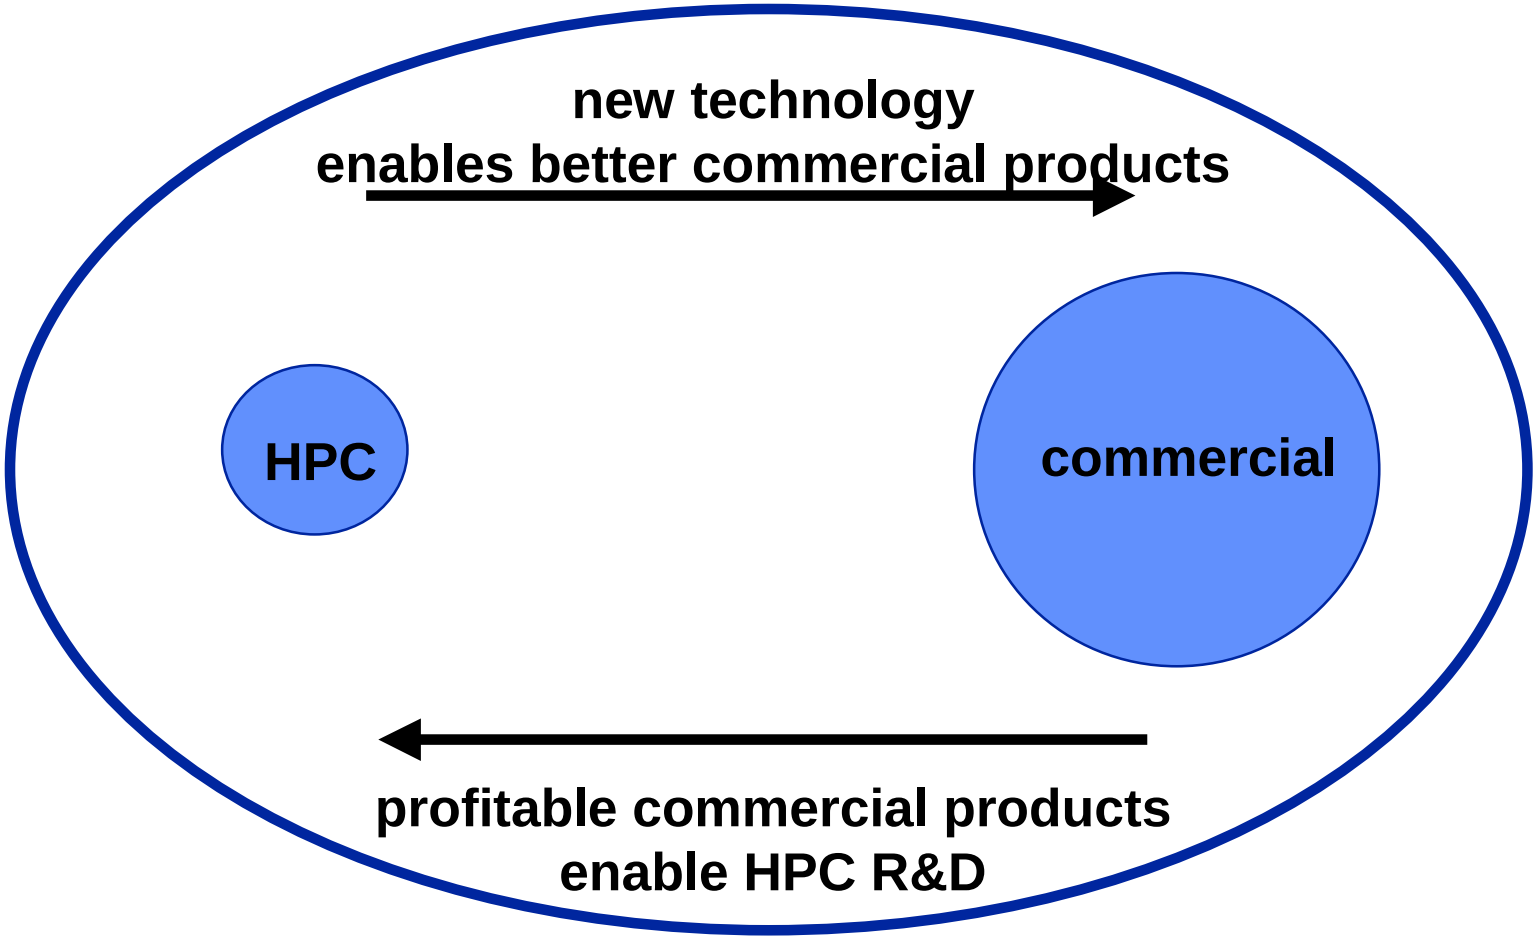

1997: The New Supercomputer Center

In 1995-97 DOD, DOE, and NSF competitively reexamined the role of centers

- **Rapidly changing technology**
- **Better local facilities everywhere**
- **Growth of computational approaches in all disciplines**

New Model: Intellectual Services + a Major Facility

New algorithms and strategies developed in medium and long-term collaborations with scientific user community

The Center is the working interface between computer science and physical science

Necessary but not sufficient

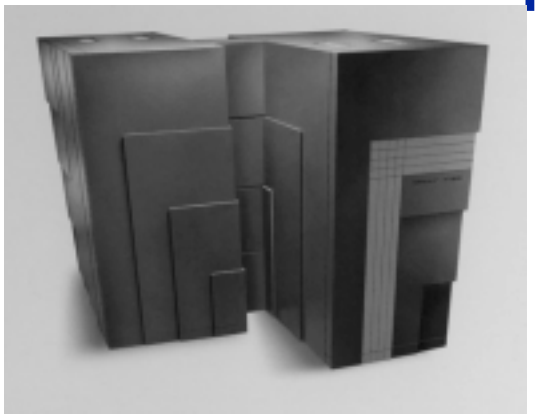

1997: The New Supercomputer Center

NSF (National Science Foundation)

Re-competition of four national centers (Cornell, NCSA, Pittsburgh, San Diego).

Two winners (NCSA, San Diego) announced in March ‘97.

New alliances for superocmputing (NPACI)

DOE (Department of Energy) - Energy Research

Competition for NERSC between Livermore and Berkeley.

Center moved to Berkeley in 1996, at 20% reduction in budget.

1997: The New Supercomputer Center

NASA (National Aeronautics and Space Administration)

Consolidation of all superocmputer operations into one activity, COSMO.

DOD (Department ofDefense)

Defense modernization program; consolidation of all superocmputing into four production centers (WES and NAVO Missisippi, Wright-Patterson, Ohio, and ARL, Maryland);

as well as two research centers (Maui and Minnesota).

1997: The New Supercomputer Center

Smaller number of larger centers

Intellectual leverage from universities

All are based on highly parallel technology

Cray T3E-900 at NERSC

Largest unclassified supercomputer in the U.S.

Scalable performance up to 512 nodes demonstrated even on “hard” applications such as climate and materials

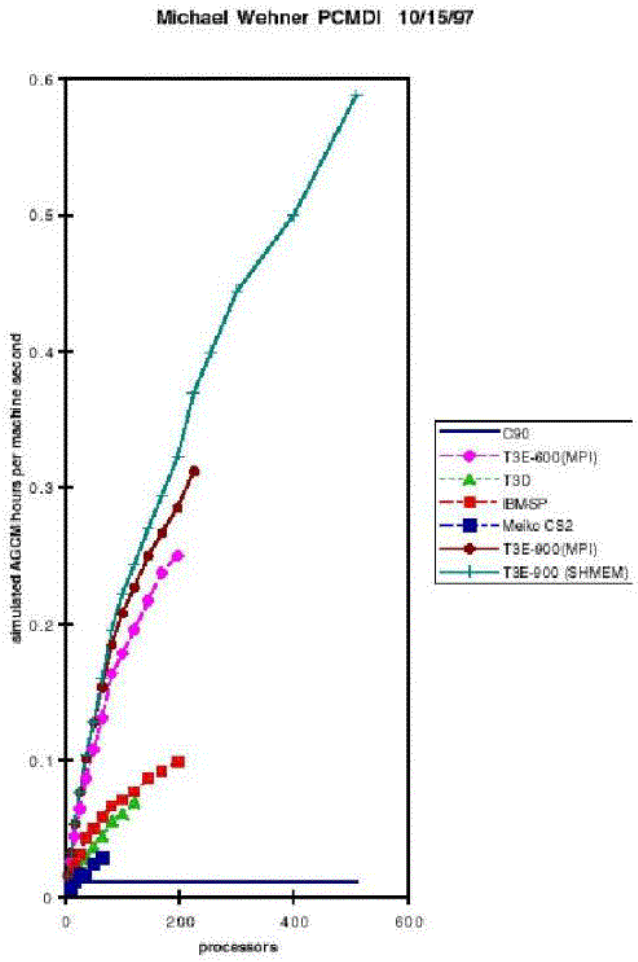

Source : Michael Wehner, PCMDI, LLNL
http://www-pcmdi.llnl.gov/wehner/mikefig.html

Overview

1992-1997 **1997** **1997-2002**

technology
markets/business
government

future non-linear events

1992-1997: Non-linear event

Internet Growth

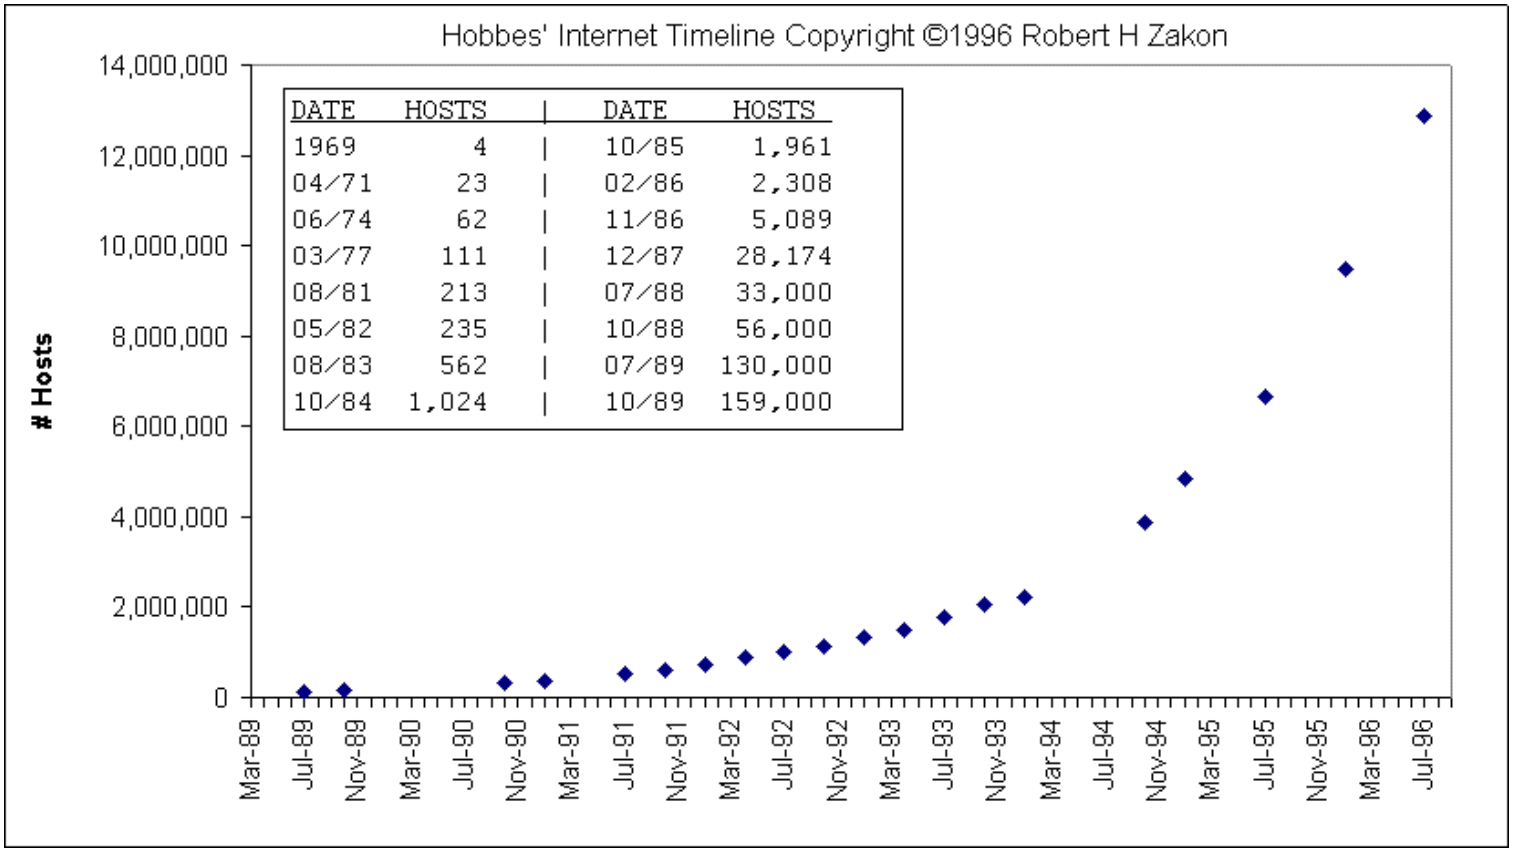

DATE	HOSTS	DATE	HOSTS
1969	4	10/85	1,961
04/71	23	02/86	2,308
06/74	62	11/86	5,089
03/77	111	12/87	28,174
08/81	213	07/88	33,000
05/82	235	10/88	56,000
08/83	562	07/89	130,000
10/84	1,024	10/89	159,000

Source http://info.isoc.org/guest/zakon/Internet/History/HIT.html
copied with permission

Internet Growth

Growth a combination of factors

- **WWW, browsers**
- **inexpensive multi-media PCs for contents**
- **NOT network bandwidth**

Impact on HPC

- **no more hot start-ups, venture capital**
- **no more new federal mega programs**
- **no longer attractive for recent graduates**

Internet2

Internet2 **http://www.internet2.edu/index.html**

- **cooperative effort of 109 member univerisities**
- **focus on reearch and higher education**
- **goals similar to NGI**
- **realtionship to NGI being defined**

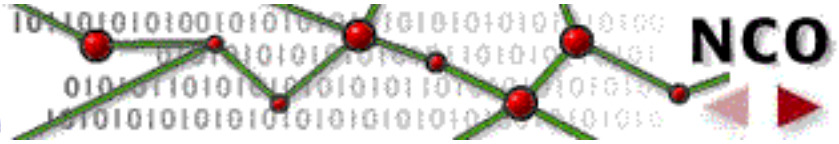

Next Generation Internet (NGI) http://www.ngi.gov/

- **$100M/year presidential initiative**

•"21st Century environment for business, education, culture, and entertainment. Sight, sound, and even touch will be integrated through powerful computers, displays, and networks. People will use this environment to shop, bank, study, entertain, work, and visit with each other."

- high performance networks (100 sites at 100 times speed, 10 sites at 1000 times speed)

- advanced network services (QoS, security, collaboratories ...)

- new applications (health care, education, research..)

1997-2002: Government

ASCI - Accelerated Strategic Computing Initiative

http://www.llnl.gov/asci/

1996 comprehensive testban on nuclear weapons signed;

shift from nuclear test-based methods to computational-based methods of ensuring the safety, reliability, and performance of nuclear weapons stockpile

create predictive simulation and virtual prototyping capabilities based on advanced weapon codes

accelerate the development of high-performance computing far beyond what might be achieved in the absence of a focused initiative.

ASCI (cont.)

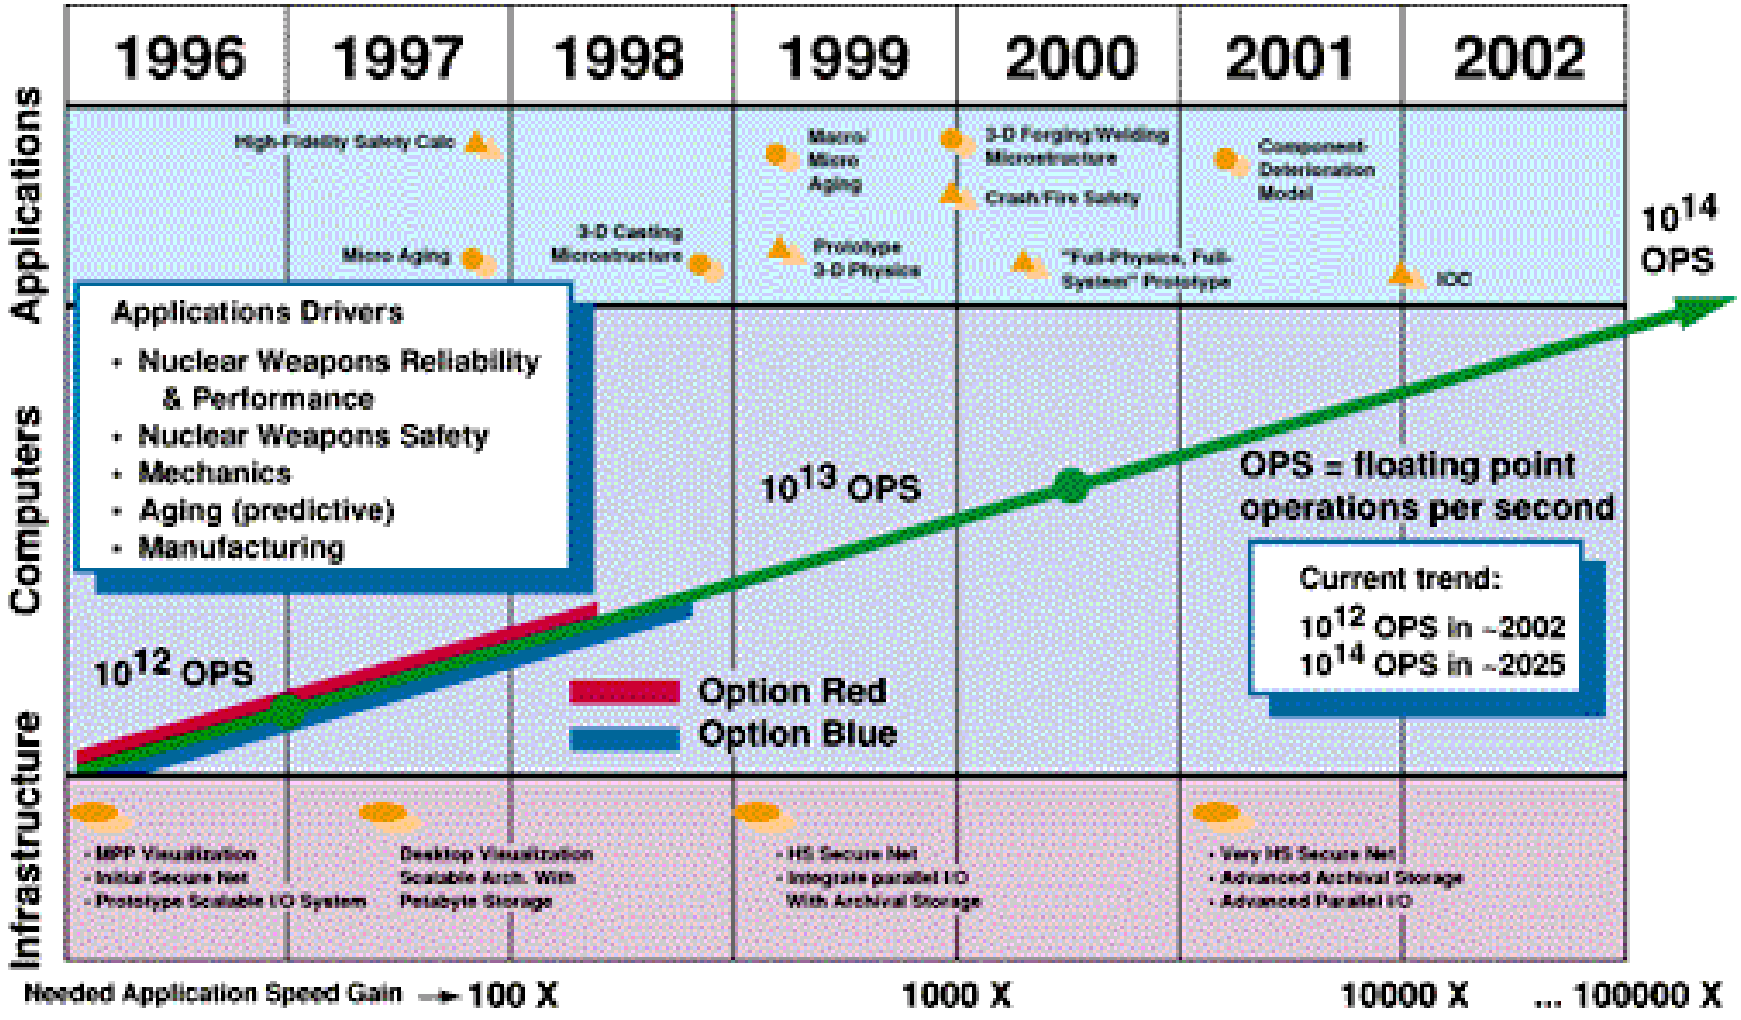
1996
1997
1998
1999
2000
2001
2002
Applications
Computers
Infrastructure
High-Fidelity Safety Calc
Micro Aging
3-D Casting Microstructure
Macro/ Micro Aging
Prototype 3-D Physics
3-D Forging/Welding Microstructure
Crash/Fire Safety
"Full-Physics, Full-System" Prototype
Component Deterioration Model
IDC
Applications Drivers
• Nuclear Weapons Reliability & Performance
• Nuclear Weapons Safety
• Mechanics
• Aging (predictive)
• Manufacturing
10^{14} OPS
10^{13} OPS
10^{12} OPS
OPS = floating point operations per second
Current trend:
10^{12} OPS in ~2002
10^{14} OPS in ~2025
Option Red
Option Blue
• MPP Visualization
• Initial Secure Net
• Prototype Scalable I/O System
Desktop Visualization
Scalable Arch. With Petabyte Storage
• HS Secure Net
• Integrate parallel I/O With Archival Storage
• Very HS Secure Net
• Advanced Archival Storage
• Advanced Parallel I/O
Needed Application Speed Gain
100 X
1000 X
10000 X
... 100000 X

ASCI Strategy

1. Seamless Management: One Program-Three Labs (LANL, LLNL, Sandia)

2.Focus on Advanced Applications Development

    Focus on 3-D, high-fidelity, full-systems weapons simulation applications; verification

3.Focus on High-End Computing

    ASCI Red at Sandia (Intel); ASCI Blue Mountain at LANL (SGI); ASCI Blue Pacific (IBM) at LLNL

4.Create Problem-Solving Environments

5.Encourage Strategic Alliances and Collaborations

    about five university research centers at $5M/year for 5 years to be announced shortly

ASCI in 2002 (opinion)

... will have made profound impact on the acceptability of computational modeling in science and engineering

... will turn computation from service to science in the lab culture

... will probably NOT reach the 100 Tflops (LINPACK) goal in 2002 (ASCI has a large budget, but this is not sufficient to change technology direction)

1997 - 2002: Technology

Continued growth of microprocessor based performance assured

- **Paper by Carver Mead sees no immediate limit**
- **~1 Gflop/s microprocessors by 1997**
- **~4 Gflop/s microprocessors by 2000**

Consequently:

- **no more custom architectures**
- **no massive parallelism**
- **moderate microprocessor based parallelism will be the norm**

1997 - 2002: Technology

Microprocessor Performance Will Continue to Grow

With superscalar, out order execution, branch predicition, speculative execution, large caches, increasing clock rates

Price Performance Will Also Improve

Custom architectures increasingly not viable

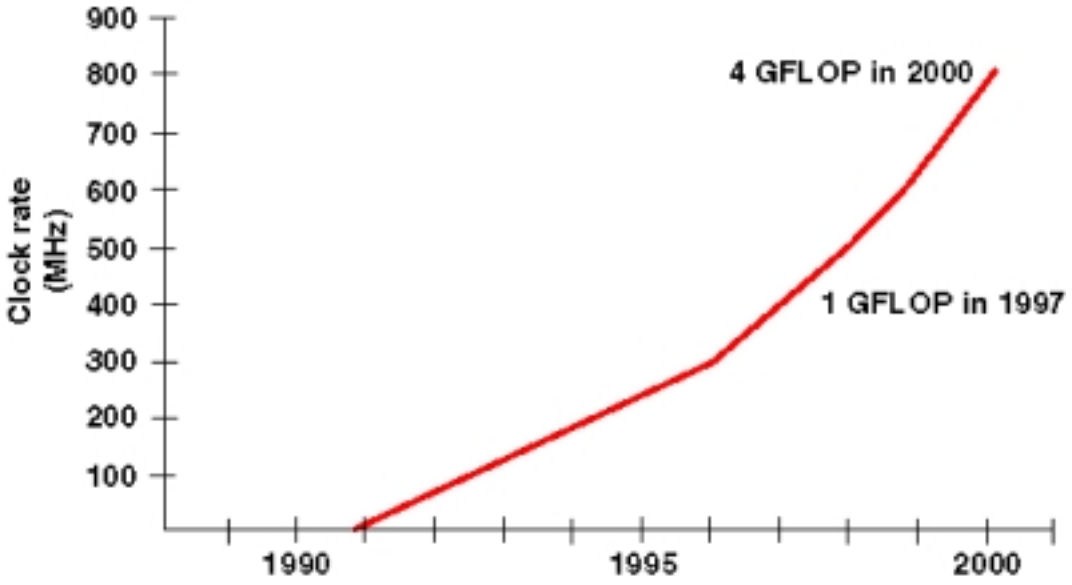

TOP500 Data 1993 - 96

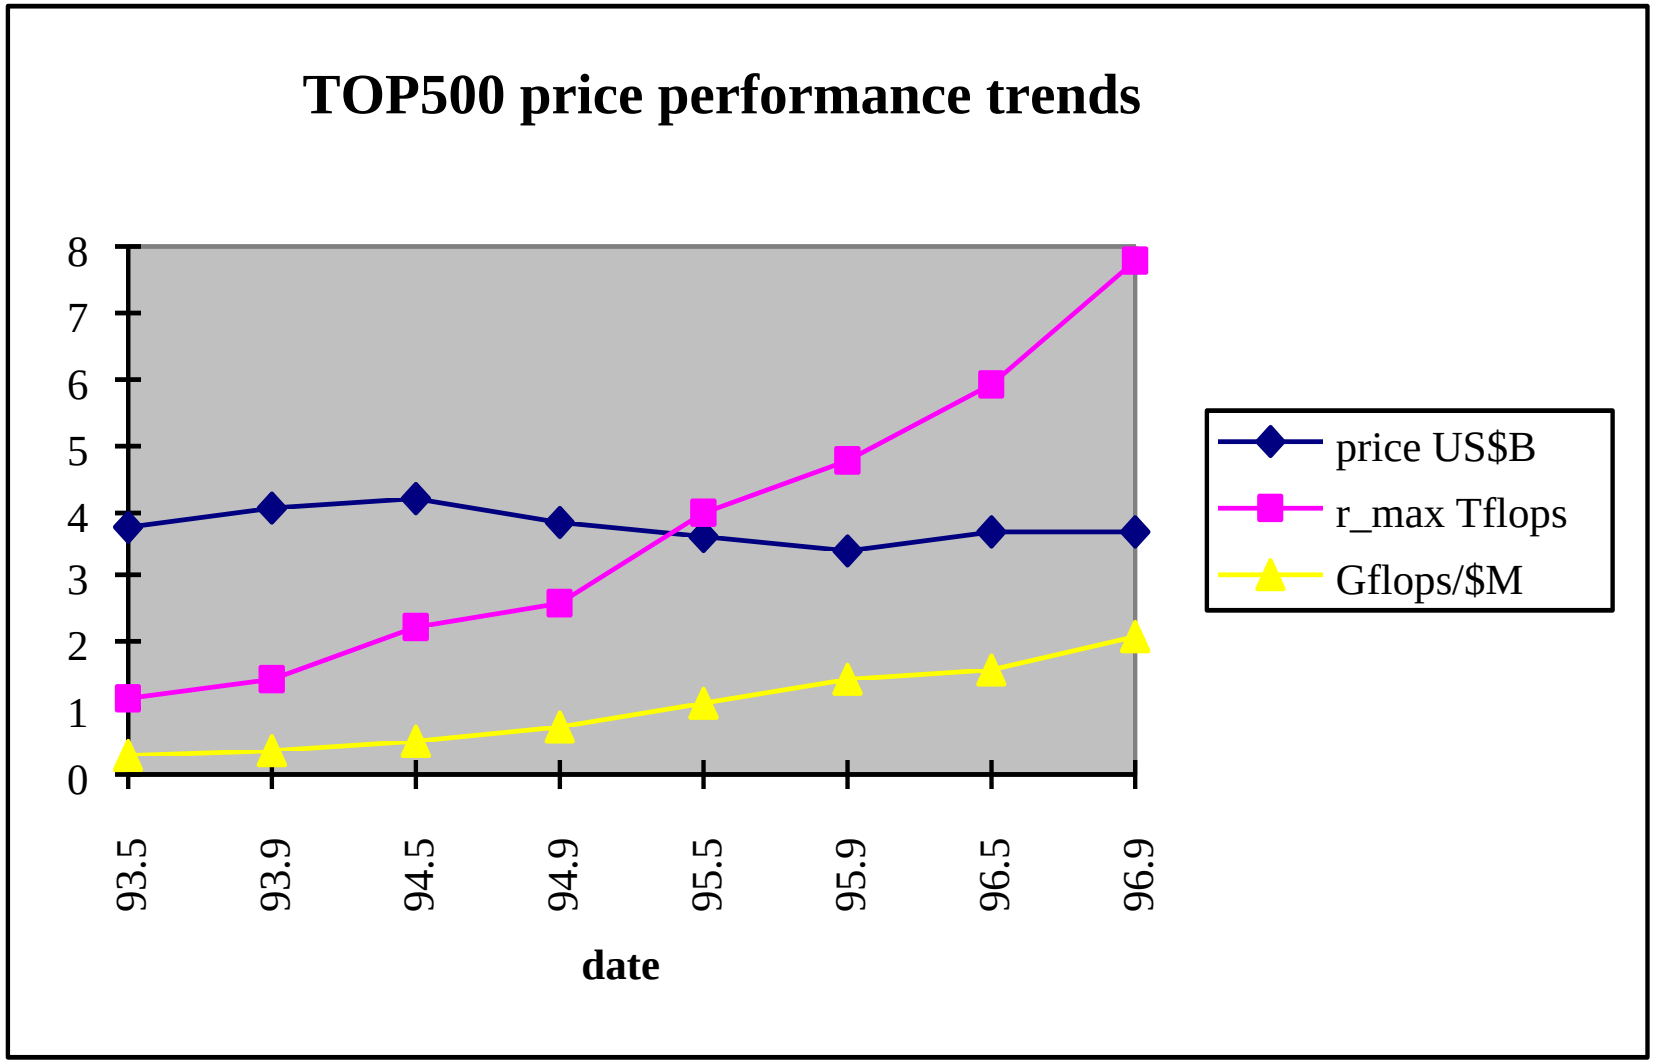

TOP500 Performance Prediction for 2002

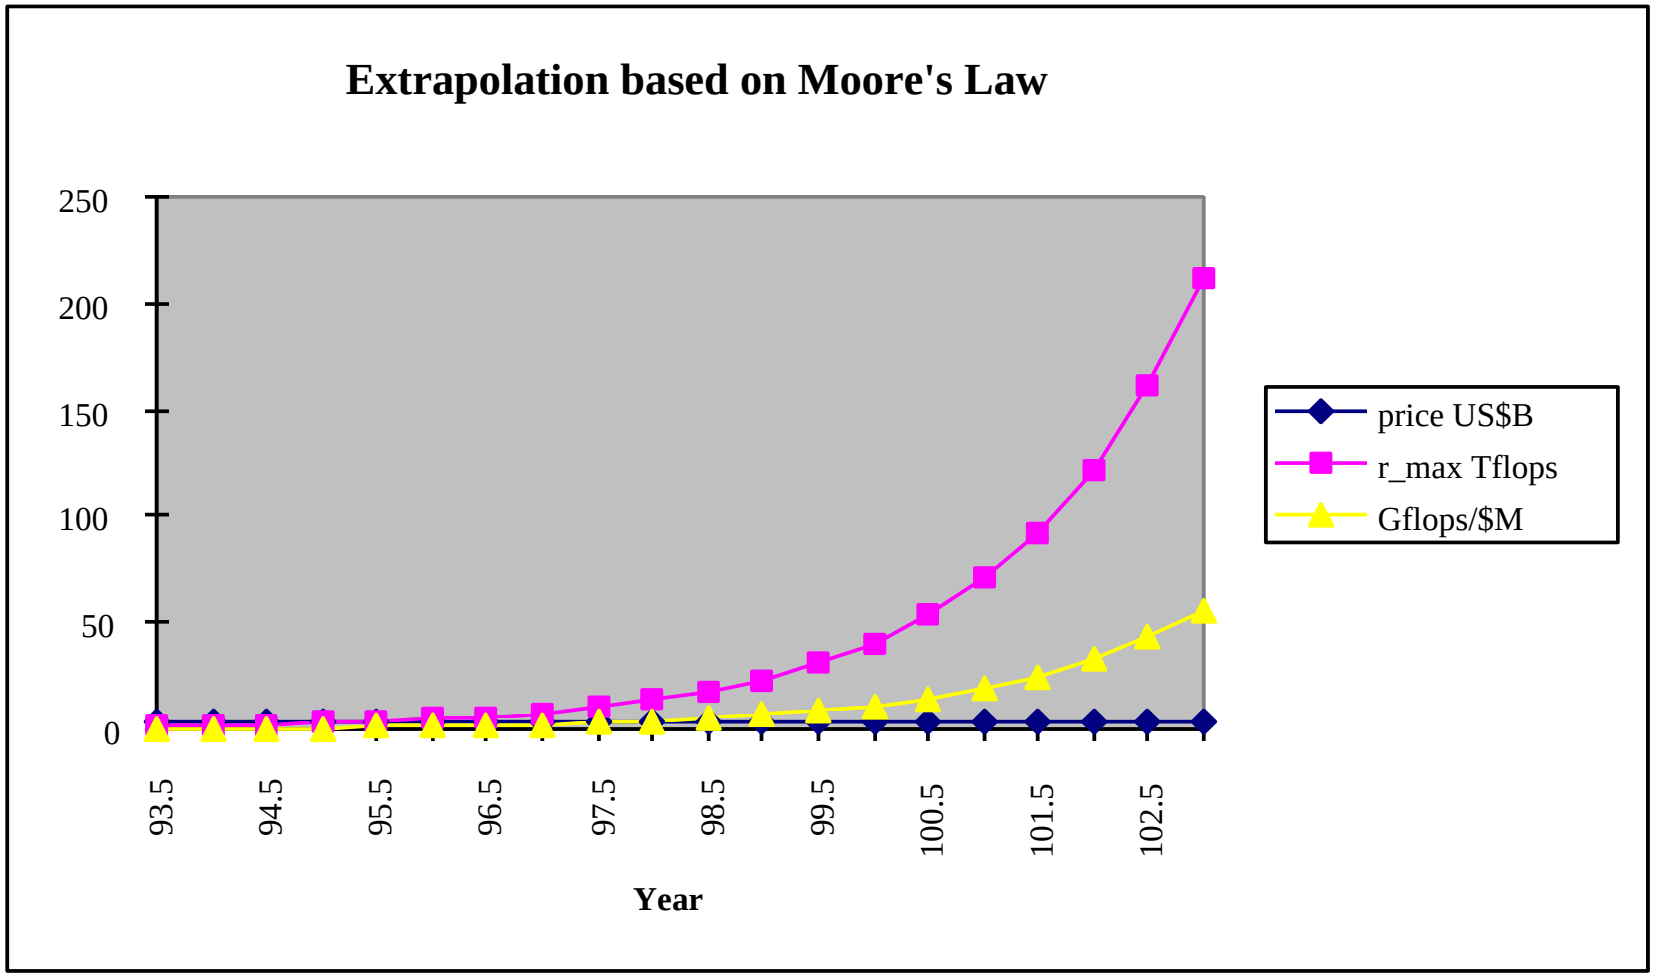

TOP500 Performance Prediction for 2002

	6/1997	6/2000	11/2002	
R_max for#1	1.0	5.0	22	Tflops
total R_max	10.2	53.4	211	Tflops
perf. of $1M system	2.7	14.2	56.3	Gflops

This is based on processor performance and increase in parallelism.

1997 - 2002: Market Issues

From vertical to horizontal companies - the Compaq model of High Performance Computing

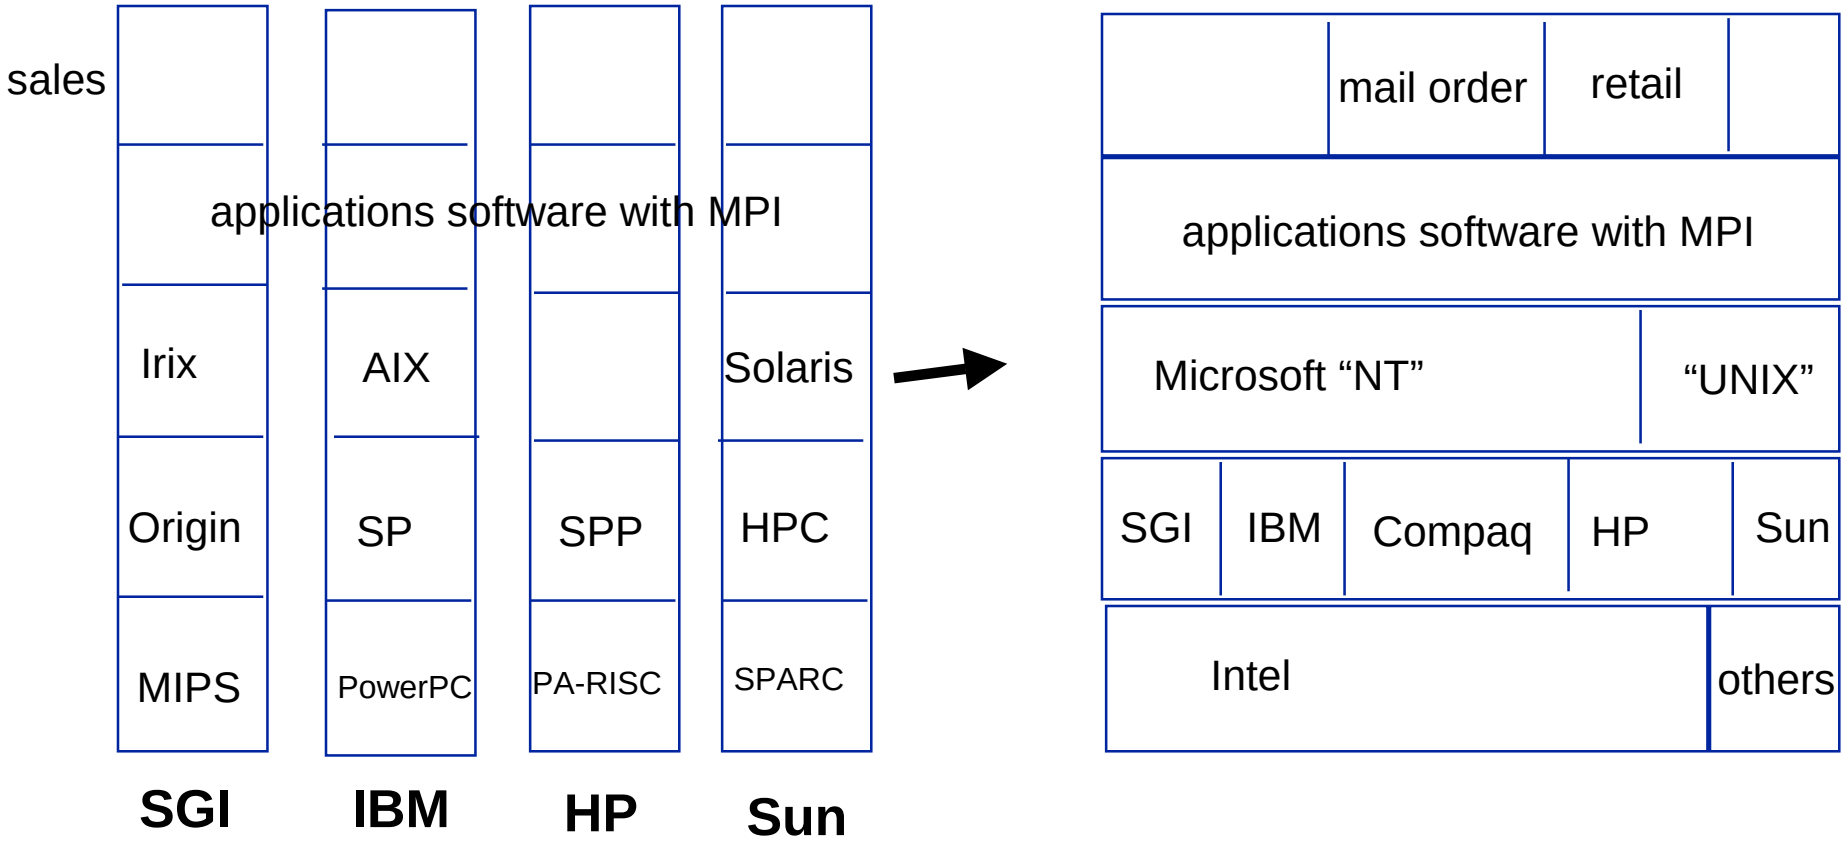

1997-2002: Market Issues

- Compaq's acquisition of DEC is just the first step.
- DEC transformed from vertical to horizontal in less than one year.
- Business transition will be more fundamental than previous technology transition.
- Tremendous impact on HPC community - no more business as usual (e.g. how do we procure machines)
- Extremely difficult to pick winner
- Tumultous transition may create niches for boutique companies such as Tera to survive

http://www.businessweek.com/1997/31/970804.htm

1997 - 2002: Ubiquitous Computational Modeling

Commodity consumer products

Example:

MOTOROLA, Pager Division, Boynton Beach, Florida

Applications: Radioss/Parallel Solids

ABAQUS Standard/Explicit

Alias - Render Industrial Designs

EFMASS, MDS, from H.P., MCSPICE

System: **8 CPU POWER CHALLENGE**

2 GB Memory, 40GB Disk

Problem: **Pager Case**

- **Battery Containment**
- **Electronics Integrity**
- **Display Life**

1997 - 2002: Ubiquitous Computational Modeling

1985	1990	1995
specialized hardware	specialized hardware	commodity hardware
Cray X-MP	Cray Y-MP	POWER CHALLENGE XL
nuclear weapons lab.	industrial company	industrial company
	unique control resource	decentralized divisonal resource
unique multimillion $ product	expensive consumer product	mass consumer product
(weapons impact)	$10K	$1.99
	(car crash)	(pager/cellular phone)

The Economic Limit: Moore's Second Law

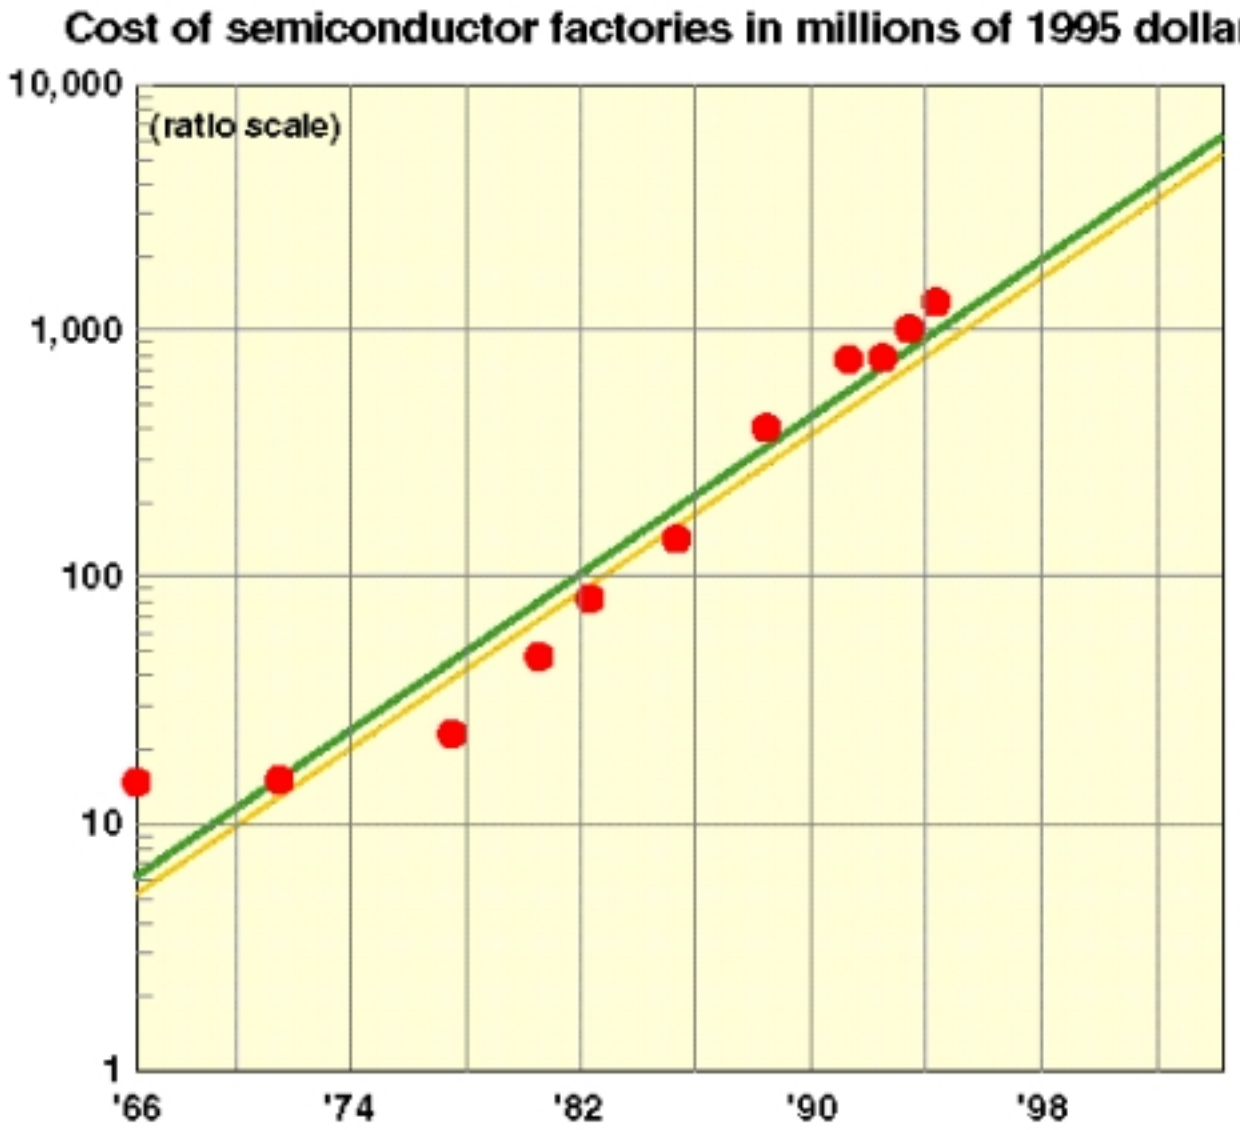

Source: Forbes Magazine

1997-2002: New Technology

IRAM - Intelligent Random Access Memory

http://iram.cs.berkeley.edu/

- over the next decade processors and memory will be merged onto a single chip

- remove processor-memory performance gap

- ideal building-block for parallel processing

- amortize the costs of fabrication lines

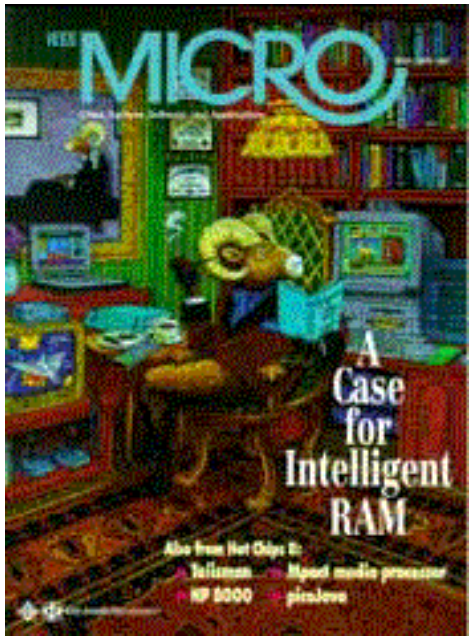

1997-2002: New Technology

Petaflops Initiative - series of workshops

not a “program” by any agency

four approaches to Petaflops systems:

- **conventional DSM silicon architecture**
- **superconducting design**
- **processor in memory (~ IRAM)**
- **special purpose architectures**

1997-2002: New Technology

The software challenge: overcoming the MPI barrier

- **MPI created finally a standard for applications development in the HPC community**
- **standards are always a barrier to further development**
- **the MPI standard is a least common denominator building on mid 80ies technology**

Summary

- After a revolutionary technology transition the years 1997-2002 promise rapid growth in HPC
- MPPs have matured and will benefit greatly from microprocessor technology; this will be the “golden” MPP era
- The HPC industry will undergo a fundamental transition to a horizontal model, bringing major uncertainties for the HPC community
- Moore’s second law will force HPC community to look seriously at massively parallel computing (after 2002) or develop new technology
- Federal agencies must anticipate these rapid changes and initiate new programs